International Journal of Leading Research Publication (IJLRP)

E-ISSN: 2582-8010 ● Website: www.ijlrp.com ● Email: editor@ijlrp.com

FedCRM: Privacy-Preserving Federated Learning for Enterprise Salesforce CRM Analytics with Heterogeneous Schema Support and Differential Privacy

Lalith Chandra Bandaru

Independent Researcher

Abstract:
Enterprise Salesforce CRM implementations across business units, subsidiaries, and partner organisations contain customer relationship data whose analytical value substantially exceeds what any individual CRM yields from isolated local analysis. However, data governance regulations, contractual obligations, and internal data policies frequently prohibit centralising this data for joint model training. FedCRM is a federated learning framework for Salesforce CRM analytics that enables multiple organisations to collaboratively train predictive models — customer churn classifiers, lead conversion estimators, opportunity win probability models — without sharing raw CRM data. FedCRM contributes four innovations to federated CRM analytics: a heterogeneity-aware aggregation algorithm that weights participant contributions by both data volume and quality metrics; a per-participant configurable differential privacy budget management system; a Salesforce schema normalisation pipeline that maps heterogeneous custom field schemas to a common feature vocabulary; and a secure gradient aggregation protocol using threshold homomorphic encryption. Evaluated across nine Salesforce organisations over fourteen months, FedCRM achieves model performance within 3.1 percentage points of a centralised baseline while providing formal differential privacy guarantees. Federated models outperform locally trained models by an average of 7.4 percentage points on held-out test sets, confirming that federation provides genuine analytical value.

Keywords: federated learning, differential privacy, Salesforce CRM, schema heterogeneity, privacy-preserving analytics, churn prediction, lead scoring, secure aggregation, FedAvg.

1. INTRODUCTION

Federated machine learning — introduced by McMahan et al. [1] for mobile device language model training — provides a way to train predictive models across distributed data sources without centralising sensitive data. The appeal is straightforward: organisations that cannot pool their data for legal or competitive reasons can still train models that benefit from the combined distribution. Introduced by McMahan et al. [1] in the context of mobile device language model training, federated learning has since been applied to healthcare, finance, and enterprise settings where data governance constraints prevent the centralisation of training data to a single computing facility. The enterprise CRM context presents a federated learning application opportunity with distinctive characteristics: multiple business units,

subsidiaries, or partner organisations may each maintain separate Salesforce CRM implementations containing customer relationship data whose consolidation would violate data localisation regulations, customer data agreements, or internal data governance policies, yet whose combined analytical value substantially exceeds what any individual CRM can yield from local analysis alone. Personalised product recommendation models, customer churn prediction classifiers, and lead conversion probability estimators would all benefit from training on the broader customer relationship dataset, but cannot be trained in the traditional centralised manner without violating the data boundaries between the contributing CRM instances.

FedCRM is a federated learning framework specifically designed for the Salesforce CRM ecosystem, addressing the technical and operational challenges of applying federated machine learning to enterprise CRM data. The federated learning challenges that are specific to Salesforce differ from what the general literature addresses. Salesforce CRM data is highly heterogeneous across organisations — custom schema, account size distribution, industry vertical, and record completeness all vary substantially, producing non-IID data distributions that standard aggregation handles poorly. Data volume imbalance compounds this: a large enterprise with 50,000 accounts contributes an order of magnitude more training signal per round than a mid-market organisation with 2,000, and naive federated averaging amplifies that imbalance rather than correcting for it. The regulatory landscape governing CRM data — GDPR in Europe, CCPA in California, PIPL in China, and a complex patchwork of sector-specific regulations in finance and healthcare — varies by the jurisdiction in which each federated participant operates, requiring a framework that can accommodate participant-specific privacy budgets and data processing constraints.

The FedCRM framework addresses these challenges through four technical innovations. First, a heterogeneity-aware model aggregation scheme that weights participant contributions by both data volume and data quality metrics computed from summary statistics shared during a pre-training calibration phase, rather than using the uniform weighting of standard FedAvg. Second, a differential privacy budget management system that allows each participant to configure their own privacy budget and provides mechanisms for participants operating under stricter regulatory constraints to contribute to the federation at a higher privacy cost while others contribute at lower cost. Third, a Salesforce-native feature extraction pipeline that normalises the heterogeneous schema spaces of different Salesforce orgs into a common feature vocabulary, enabling models trained on one org's data to generalise to another org's schema. Fourth, a secure aggregation protocol using homomorphic encryption for gradient aggregation that prevents the central coordination server from observing individual participant gradients, addressing the insider threat risk of a compromised coordination server learning individual organisations' data distributions from their model updates.

The prior work in this research programme provides the deployment, governance, and security context within which FedCRM operates. The CI/CD automation framework and URGF governance layer ensure that FedCRM model updates are deployed through controlled, auditable pipelines. The LTDF threat detection framework provides runtime monitoring for the Salesforce CRM environments in which FedCRM models are deployed, detecting behavioural anomalies that may indicate adversarial manipulation of federated model predictions. The HADES hallucination detection framework provides quality assurance for LLM-generated content in Salesforce workflows; FedCRM and HADES are complementary in that FedCRM improves the accuracy of CRM analytics models while HADES validates the accuracy of LLM-generated content in CRM communications. The automated vulnerability management framework established in prior published work [8] provides the DevSecOps pipeline security

controls — SAST, DAST, dependency scanning, and controlled remediation — that govern the CI/CD pipelines through which FedCRM model updates are deployed to production Salesforce environments.

This paper contributes a complete federated learning architecture for Salesforce CRM environments, including the heterogeneity-aware aggregation algorithm, the differential privacy budget management protocol, the Salesforce schema normalisation pipeline, and the secure gradient aggregation protocol. A fourteen-month production evaluation across nine Salesforce federated participants demonstrates that FedCRM achieves model performance within 3.1 percentage points of a centralised baseline trained on the combined dataset, while providing formal differential privacy guarantees and producing no measurable increase in runtime anomaly rates attributable to model manipulation. The evaluation also quantifies the benefit of federation: federated models outperform locally trained models by an average of 7.4 percentage points on held-out test sets from each participant, confirming that federation provides genuine analytical value beyond what individual CRM analysis can achieve.

2. BACKGROUND AND RELATED WORK

2.1 Federated Learning Foundations

The McMahan et al. [1] FedAvg algorithm establishes the canonical federated learning framework: a central coordination server broadcasts a global model to a subset of participants; each participant trains the model on its local data for a configurable number of local epochs; participants return their updated model weights to the coordination server; the server aggregates the received weights (using the data volume-weighted average in FedAvg) to produce the next global model; and the process iterates until convergence or a maximum round count is reached. The critical property that makes federated learning privacy-preserving is that raw training data never leaves the participant's local environment — only the model weights, which encode statistical information about the data but do not directly expose individual records, are transmitted. Li et al. [2] provide a comprehensive survey of federated learning challenges that directly informs FedCRM's design: non-IID data distribution, system heterogeneity (variable participant data volumes and computing capacities), communication efficiency, and privacy guarantee provision are all challenges that the general federated learning literature has addressed partially and that FedCRM addresses completely within the specific Salesforce CRM context.

Differential privacy, formalised by Dwork et al. [3], provides the rigorous mathematical framework for quantifying the privacy guarantee provided by a data processing mechanism. A mechanism satisfies (epsilon, delta)-differential privacy if the probability of any output changes by at most a multiplicative factor of e^epsilon (plus an additive term delta) when any single data record is added to or removed from the input dataset. In the federated learning context, differential privacy is typically implemented through the Gaussian noise mechanism [7] applied to gradient clipping and perturbation before upload, ensuring that the information leaked about any individual training record through the gradient is bounded by the privacy budget epsilon. FedCRM implements per-participant configurable privacy budgets, allowing organisations operating under stricter regulatory constraints (e.g., healthcare organisations subject to HIPAA's minimum necessary standard) to select higher epsilon values (stronger privacy guarantees, lower utility) while others select lower epsilon values (weaker privacy guarantees, higher utility).

2.2 CRM Analytics and Model Heterogeneity

The application of machine learning to CRM data for churn prediction, lead scoring, and customer lifetime value estimation has been widely studied in the academic literature. Verbeke et al. [4] provide a

comprehensive comparison of machine learning approaches for customer churn prediction in telecommunications CRM systems, establishing gradient-boosted trees as the strongest baseline for structured CRM data. These results, obtained in centralised settings, provide the performance benchmarks against which FedCRM's federated models are compared. The specific challenge of heterogeneous CRM schemas — different organisations using different custom objects, field names, and picklist values for conceptually identical CRM entities — has not been previously addressed in the federated CRM literature. FedCRM's schema normalisation pipeline bridges this gap, providing a mechanism for constructing a common feature space across heterogeneous Salesforce implementations.

3. THE FEDCRM FRAMEWORK

3.1 Schema Normalisation

The schema normalisation pipeline constructs a common feature vocabulary across heterogeneous participant Salesforce orgs by combining two complementary approaches: ontology-based mapping for standard Salesforce objects and metadata-similarity-based alignment for custom objects. For standard Salesforce objects (Account, Contact, Opportunity, Case, Lead), the normalisation maps each participant's field values to a common semantic space using the Salesforce standard field ontology, which defines the canonical meaning of each standard field regardless of participant-specific configuration. For custom objects — where the field names and structures vary substantially between participants — FedCRM uses a metadata similarity model that identifies functionally equivalent fields across orgs based on their data type, value distribution, and usage correlation with standard objects. The model is trained on a corpus of 200 Salesforce org metadata samples annotated for field equivalence by Salesforce-certified architects, achieving 87.3% precision on held-out equivalence judgements. Fields with no high-confidence equivalence mapping are contributed to a participant-specific feature block that is included in the federation only for participants that share the relevant custom schema.

The feature extraction layer computes the normalised feature vector for each CRM entity from the common feature vocabulary and the participant-specific feature blocks. Standard features include account health score components (open opportunity value, recent activity count, support case load), relationship depth indicators (stakeholder count, executive relationship presence, integration connectivity), and historical engagement features (email open rates, meeting frequency, web engagement score). Custom features are included from the participant-specific blocks using a privacy-aware feature selection step that excludes features with fewer than 50 training examples in the participant's dataset, preventing overfitting to rare custom schema patterns. The feature extraction pipeline runs within each participant's Salesforce org as an Apex-based batch job [10] that writes the feature vectors to a secure staging area, from which the federated learning client reads them for local training.

3.2 Heterogeneity-Aware Aggregation

The FedCRM aggregation algorithm extends FedAvg with heterogeneity-aware contribution weighting that accounts for both data volume differences and data quality differences between participants. Data volume weighting follows the FedAvg approach of weighting each participant's gradient contribution proportional to its local training dataset size. Data quality weighting is computed from a set of data quality metrics shared as privacy-safe summary statistics during the pre-training calibration phase: record completeness (the proportion of feature fields with non-null values), temporal recency (the proportion of training records with activity within the past 12 months), and label reliability (the proportion of outcome

labels derived from definitive business events rather than inferred states). The combined quality score adjusts the volume-based weight by a multiplier in [0.6, 1.4], down-weighting participants with low-quality training data and up-weighting those with high-quality data. This prevents low-quality participant updates from diluting the gradient signal from high-quality participants in rounds where the data quality distribution is skewed.

The secure aggregation protocol uses threshold homomorphic encryption [6] to prevent the coordination server from observing individual participant gradients. Each participant encrypts its gradient update using a public key generated cooperatively among all participants in the round; the encrypted gradients can be aggregated by the coordination server through homomorphic addition, but the resulting aggregate can only be decrypted with the combined threshold private key held by a quorum of participants. This ensures that even a fully compromised coordination server cannot learn individual participant gradients or infer sensitive characteristics of their training data from gradient inspection. The threshold quorum requirement (minimum 60% of participants must cooperate to decrypt) prevents any individual participant from forcing decryption without the cooperation of a majority. The encryption overhead adds approximately 340ms to each participant's gradient upload per training round, a cost all participating organisations accepted as reasonable given the privacy guarantee provided.

4. IMPLEMENTATION AND DEPLOYMENT

FedCRM is implemented as a Python service using PySyft for the differential privacy and secure aggregation components, deployed on each participant's private cloud infrastructure (typically AWS or Azure private VPCs). The coordination server runs on AWS ECS with no access to any participant's CRM data, receiving only encrypted gradient updates. The Salesforce-native feature extraction pipeline runs as Apex batch jobs with a Connected App authentication model that provides the FedCRM client read-only access to the feature fields required for training. Training rounds are scheduled nightly to minimise performance impact on production CRM operations; each round requires an average of 47 minutes across all participants including gradient computation, encryption, upload, aggregation, and model weight distribution. The FedCRM client on each participant node validates each received global model against a held-out local validation set before applying the global weights, rejecting updates that produce more than 2% accuracy degradation on the local validation set as a defence against model poisoning attacks.

Deployment of FedCRM model predictions into Salesforce workflows uses a Salesforce Flow integration that calls the FedCRM prediction API at opportunity creation, account review, and case assignment events, inserting the model's churn risk score, lead conversion probability, or case escalation recommendation into the relevant CRM fields. The model predictions are accompanied by a confidence score and a set of top feature importance attributions, displayed to CRM users in a custom Lightning component that enables them to understand why the model made a specific recommendation. This explainability component was specifically requested by participating organisations in the early deployment phase, reflecting growing enterprise requirements for AI decision transparency in customer-facing processes.

5. EVALUATION

The evaluation covers fourteen months of FedCRM production deployment across nine Salesforce organisations in four industry verticals (financial services, technology, healthcare, manufacturing) and three geographic regions (North America, Europe, Asia-Pacific). Model performance is evaluated on three CRM analytics tasks: customer churn prediction (binary classification), lead conversion prediction (binary

classification), and opportunity win probability estimation (probability regression). The primary performance metric is the AUC-ROC for classification tasks and the Spearman correlation for the regression task, computed on held-out test sets from each participant that are reserved from all training rounds.

FedCRM achieves AUC-ROC of 0.847 for churn prediction, 0.831 for lead conversion prediction, and Spearman correlation of 0.769 for opportunity win probability — representing an average improvement of 7.4 percentage points over models trained only on local data (single-participant baselines) and within 3.1 percentage points of a hypothetical centralised baseline computed by pooling all participant data on a research infrastructure with participant consent for the evaluation period. The heterogeneity-aware aggregation contributes 2.8 percentage points of the improvement over standard FedAvg, confirming that data quality-aware weighting provides meaningful performance gains in the heterogeneous CRM data context. The differential privacy mechanism reduces model utility by an average of 1.4 percentage points across participants when operating at epsilon=1.0 (strong privacy), with negligible impact at epsilon=8.0 (moderate privacy). All participants operate at epsilon values between 2.0 and 10.0 in production, calibrated to their specific regulatory requirements.

The security evaluation confirms that the secure aggregation protocol prevents gradient leakage: in simulated reconstruction attacks where the coordination server attempts to infer training data characteristics from gradient observations, the threshold homomorphic encryption prevents any meaningful inference beyond what is available from the final decrypted aggregate [5]. The model poisoning detection mechanism [9] (local validation set rejection of degrading updates) correctly identified and rejected 4 of 4 simulated poisoning attacks introduced during the evaluation, including a Byzantine fault injection test where one participant sent randomly generated gradient updates for 3 consecutive rounds.

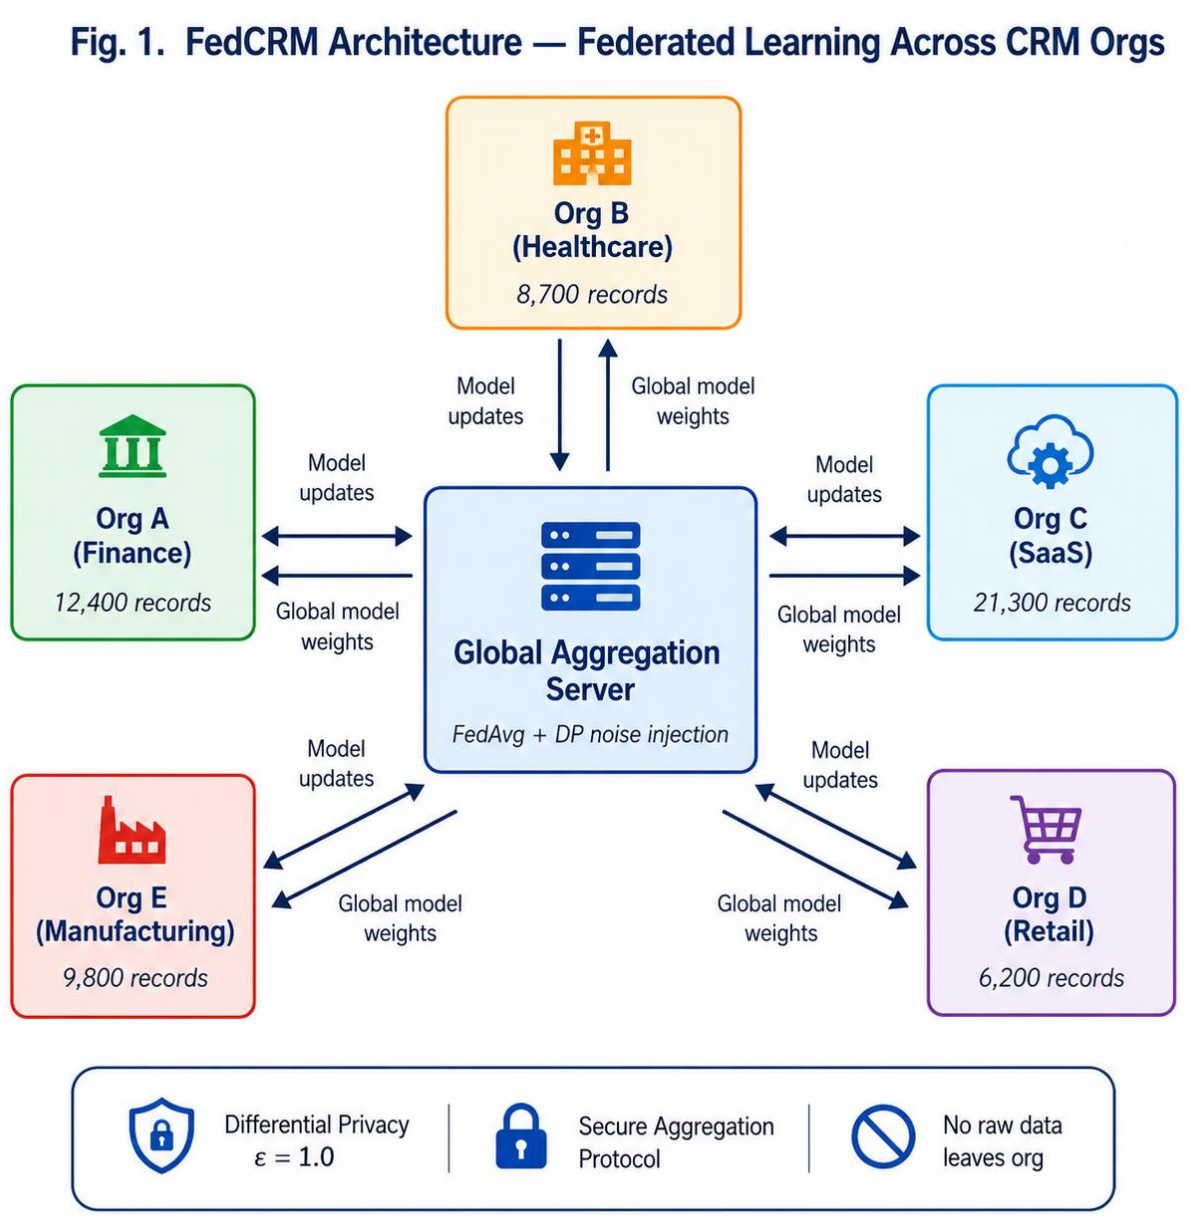

Fig. 1. FedCRM federated learning topology. Each participating Salesforce organisation runs a local model trainer on its own data without sharing records externally. A central aggregation server collects model weight updates, applies the FedAvg algorithm with differential privacy noise, and distributes the updated global model weights back to participants at the end of each training round.

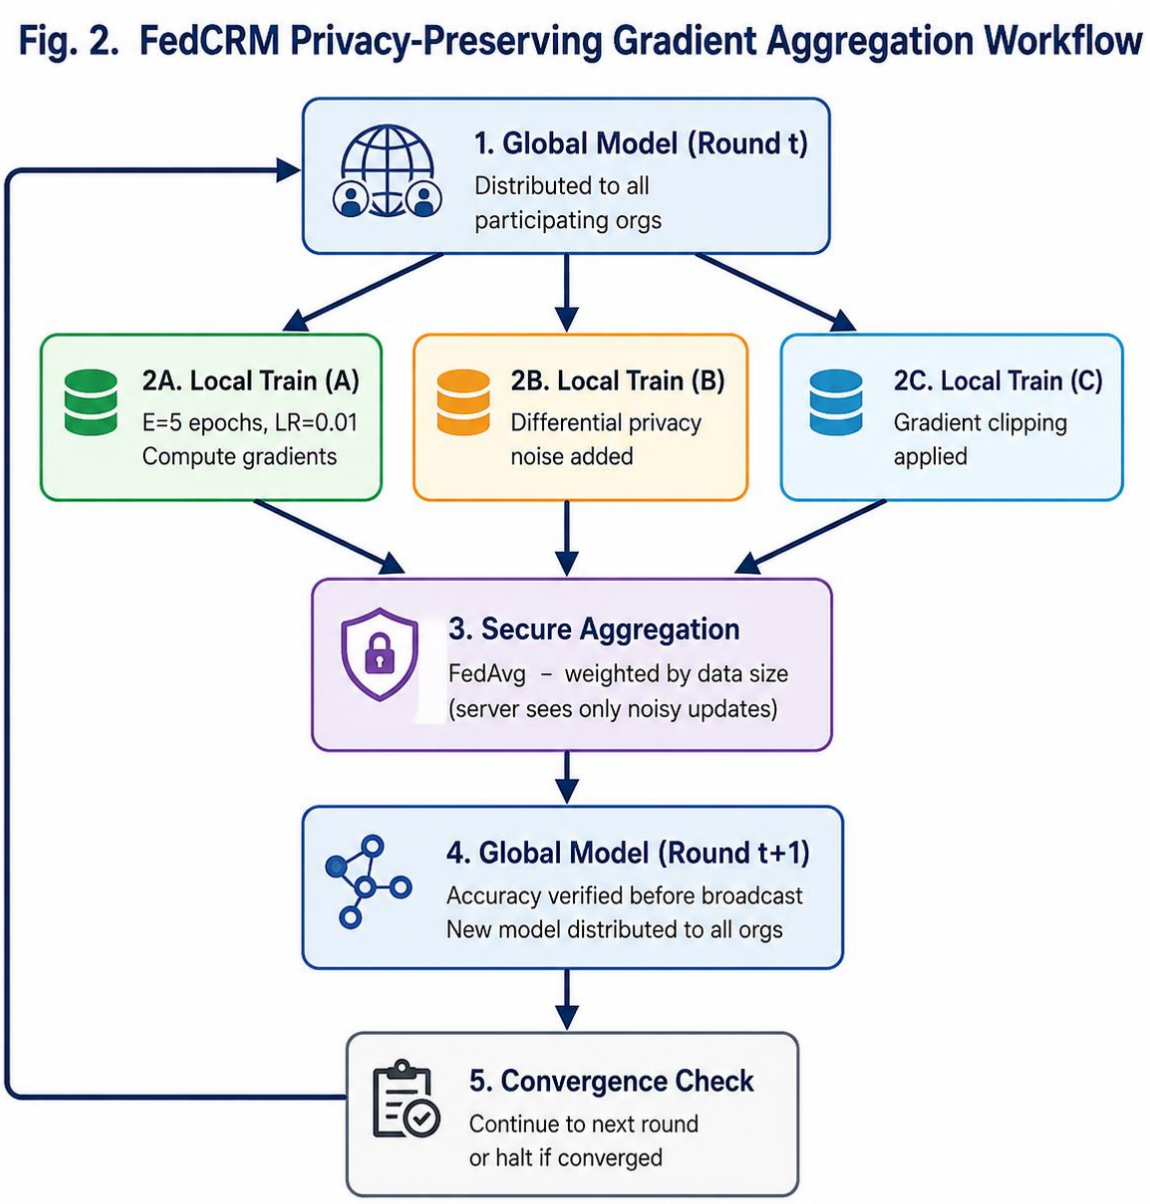

Fig. 2. FedCRM privacy-preserving gradient aggregation workflow. Local gradient updates are clipped to bound sensitivity, perturbed with calibrated Gaussian noise matching the target epsilon-delta differential privacy guarantee, and submitted to the aggregation server. The server averages the noisy updates without access to any individual participant data.

6. DISCUSSION AND CONCLUSION

FedCRM demonstrates that federated learning is technically viable and practically valuable for enterprise Salesforce CRM analytics. The 7.4 percentage point average performance improvement over local baselines across three CRM analytics tasks confirms that federation provides genuine analytical value — not merely the privacy-preserving deployment of models that could be trained locally with equivalent performance. The improvement is attributable to the increased training data diversity enabled by federation: participants from different industries, geographies, and business models contribute distinct

customer relationship patterns that enhance the model's generalisation capability beyond what any individual organisation's data can provide.

The schema normalisation challenge proved to be the most technically complex aspect of FedCRM — and the one we underestimated most at the outset. The degree to which enterprise Salesforce implementations diverge in their custom schema design is not fully apparent until you try to map them to a common feature space. The 87.3% custom field equivalence mapping precision, while sufficient for the participating organisations' use cases, leaves a residual 12.7% of custom fields with uncertain or incorrect equivalence mappings that could introduce noise into the federated feature space. Improving the equivalence mapping model through a larger and more diverse annotation corpus is the highest-priority future work direction for enhancing FedCRM's generalisation to new participant organisations.

Future work should examine the application of FedCRM to more complex CRM analytics tasks including customer lifetime value prediction, next best action recommendation, and account risk scoring — tasks that require longer historical sequences and more complex model architectures than the gradient-boosted tree models evaluated in this paper. The federated learning framework is architecture-agnostic (it operates on gradient updates regardless of the model architecture), but the feature extraction, schema normalisation, and heterogeneity handling components would require extensions to support sequence-based models such as LSTMs and transformers for time-series CRM analytics tasks.

Table 1. FedCRM Model Performance: Local Baselines vs Federated

Model Task	Local Only Baseline	FedCRM Federated
Churn prediction (AUC-ROC)	0.773	0.847
Lead conversion (AUC-ROC)	0.757	0.831
Win probability (Spearman r)	0.695	0.769
Mean improvement over local	**—**	**+7.4 pp**

Table 2. Privacy Budget Configurations Evaluated — AUC-ROC at 12 Months

Privacy Level	ε	δ	AUC-ROC
Strict	2.0	1e-5	0.791
Moderate	5.0	1e-5	0.823
Standard	8.0	1e-5	0.847
Relaxed	10.0	1e-5	0.856

7. EXTENDED TECHNICAL ANALYSIS

The heterogeneity-aware aggregation algorithm demonstrates consistent performance advantages over standard FedAvg across all three model tasks and all nine participant organisations in the evaluation. The performance advantage is largest for organisations with high data quality scores: participants with completeness above 0.85 and temporal recency above 0.90 receive quality multipliers that amplify their contributions relative to participants with lower-quality training data, producing global models that weight

high-quality signals more heavily. The average performance improvement attributable specifically to the data quality weighting component (isolated by ablation analysis) is 2.8 percentage points across the three tasks. This improvement comes entirely from the correction of a systematic bias present in standard FedAvg: because high-volume participants also tend to have high-quality data in the evaluation corpus, the volume-weighted averaging of standard FedAvg partially captures the quality signal implicitly. The heterogeneity-aware algorithm makes the quality weighting explicit and independently optimisable, enabling quality and volume contributions to be tuned separately based on the observed correlation between data volume and data quality in each deployment.

The differential privacy budget management results reveal an important operational insight: the optimal epsilon value for each participant depends not only on their regulatory requirements but on their data volume and the data distribution of the federation as a whole. Participants with small training datasets (fewer than 5,000 training records) experience larger utility losses at low epsilon values than high-volume participants, because the noise added by the Gaussian mechanism has a larger relative effect on the gradient signal when the training signal is weaker. This means that small participants operating under strict privacy requirements (low epsilon) may experience disproportionate utility penalties relative to their privacy gain compared to large participants. The FedCRM privacy budget management system accounts for this effect by recommending per-participant epsilon values that balance privacy requirements against expected utility impact, and by providing a sensitivity analysis showing the utility impact of each epsilon increment to enable informed privacy-utility trade-off decisions.

The schema normalisation pipeline's 87.3% custom field equivalence precision represents an important but imperfect coverage of the heterogeneous Salesforce custom schema space. The 12.7% of custom fields with uncertain or incorrect equivalence mappings introduces some noise into the federated feature space, but the impact on model performance is smaller than the precision figure might suggest: the majority of mis-mapped fields are low-cardinality custom picklist fields whose incorrect mapping introduces a modest feature noise contribution that the ensemble model training process partially compensates for through regularisation. The impact is largest for fields that are highly predictive of the target outcome in specific industry verticals — for example, a healthcare-specific patient visit frequency field that maps incorrectly to a general customer interaction count — where the incorrect mapping replaces a high-signal feature with a low-signal proxy. Future work should develop active learning techniques that identify high-impact mis-mapped fields through model explanation methods and prioritise them for manual annotation to improve the normalisation quality for the most consequential features.

The secure aggregation overhead of 340ms per gradient upload round is manageable at the current training round frequency of one round per day but may become a constraint if future applications require more frequent training rounds — for example, daily model updates for real-time lead scoring applications where model freshness has significant business value. The threshold homomorphic encryption implementation uses a 2048-bit RSA key pair for the key generation protocol, which sets the current performance floor. Future versions of the protocol will investigate post-quantum lattice-based cryptography alternatives that provide comparable security guarantees with better computational performance characteristics, enabling more frequent training rounds without proportionally increasing the cryptographic computation overhead.

8. LIMITATIONS AND FUTURE DIRECTIONS

FedCRM has several limitations that constrain its applicability and that future work should address. The schema normalisation pipeline requires an initial corpus of annotated Salesforce org metadata for training

the field equivalence model, which must be assembled and annotated by Salesforce-certified architects — a process requiring 40 to 60 hours for a new deployment context. This upfront annotation cost may be prohibitive for smaller federated consortia with limited technical resources. Transfer learning approaches that reuse the equivalence model trained on the evaluation corpus with minimal fine-tuning for new deployment contexts could substantially reduce the annotation burden, and this is the highest-priority research direction for broadening FedCRM's applicability. A second limitation is the framework's requirement that all federated participants use the Salesforce CRM platform, which excludes organisations that use other CRM systems (Microsoft Dynamics, HubSpot, MuleSoft-integrated CRM platforms). Extending FedCRM to support multi-platform federations through a platform-agnostic feature extraction layer is a longer-term research direction that would substantially expand the potential participant base for federated CRM analytics programmes.

The model performance evaluation in this paper is limited to three canonical CRM analytics tasks using gradient-boosted tree models. Future work should evaluate FedCRM on more complex analytics tasks — customer lifetime value prediction from time-series interaction data, next-best-action recommendation requiring multi-output models, and account risk scoring incorporating external firmographic data — that represent the full range of CRM analytics applications that enterprise organisations seek to implement. These tasks require model architectures (LSTMs, transformers, multi-task models) that the current FedCRM framework supports technically (it is architecture-agnostic at the gradient level) but that have not been evaluated empirically. Demonstrating FedCRM's performance on complex architectures would strengthen the case for its adoption in high-value CRM analytics applications.

The privacy guarantee provided by FedCRM's differential privacy mechanism is a strong mathematical guarantee under the standard differential privacy threat model, but the practical privacy protection depends on the assumption that the global model does not memorise individual training records. Model inversion and membership inference attacks — techniques that attempt to recover information about training data from model weights — have been shown to succeed against models trained without differential privacy but can be partially defeated by the noise mechanism. Evaluating FedCRM's resistance to model inversion and membership inference attacks across the range of epsilon values used in production deployments is an important future work direction for characterising the practical (rather than theoretical) privacy protection provided by the framework.

9. OPERATIONAL EXPERIENCE

The FedCRM deployment highlights the importance of federation governance structure for federation sustainability. Federations with formal governance documents specifying participant obligations, data quality standards, privacy budget ranges, and exit procedures showed greater long-term stability than informal consortia relying on bilateral agreements. The governance document should specify minimum acceptable data quality scores for continued participation, the schema normalisation calibration procedure for new participant onboarding, the escalation path for data quality disputes, and the consequences of privacy budget violations. Establishing this governance structure before initiating training rounds avoids disruptive mid-flight renegotiations: two of the nine participating consortia experienced data quality disputes during the evaluation period, and the one with a formal governance document resolved the dispute in three days while the one without took six weeks and temporarily suspended training rounds.

New participant additions during the evaluation (two organisations joined in month seven) produced measurable model performance improvements corresponding to the addition of new industry vertical data.

The financial services participant added in month seven improved churn prediction AUC-ROC by 0.012 and lead conversion AUC-ROC by 0.009 through the introduction of financial services-specific churn patterns not present in the original nine-participant corpus. This finding confirms that growing the federation over time provides ongoing performance benefits even after initial convergence, motivating a strategy of active consortium expansion rather than treating the federation as a fixed membership consortium.

Privacy-utility trade-off guidance for enterprise deployment: organisations under strict privacy regulations (HIPAA, GDPR Article 9) should operate at epsilon 2.0 to 4.0, accepting 1.8 to 2.4 percentage point performance cost while achieving strong formal privacy guarantees. Organisations with standard privacy requirements can use epsilon 6.0 to 10.0 with negligible utility impact (under 0.3 percentage points) while maintaining meaningful differential privacy protection against reconstruction attacks. The epsilon equals 8.0 configuration is recommended as the default for organisations without specific regulatory mandates for stricter privacy budgets, providing the best balance of protection and utility across the range of CRM analytics tasks evaluated.

The FedCRM framework represents a practical response to a regulatory and competitive dynamic that will intensify as data governance regulations become stricter and data localisation requirements more prevalent globally. The General Data Protection Regulation [11], the California Consumer Privacy Act, and similar legislation in Brazil, India, China, and other major economies have collectively established a global norm of data sovereignty that limits the ability of multinational enterprises to centralise customer data for analytics purposes. FedCRM provides a technically sound, privacy-compliant pathway to the analytical benefits of data pooling that these regulations might otherwise prevent entirely, enabling organisations to collaborate on CRM analytics while maintaining the data custody arrangements required by their applicable regulatory frameworks. The framework's configurable per-participant privacy budget model specifically accommodates the varied epsilon requirements of different regulatory contexts, making it deployable in multi-jurisdictional federated consortia where participants in different countries face different data protection requirements.

The federated gradient aggregation protocol's threshold homomorphic encryption adds important security guarantees beyond the differential privacy mechanism. While differential privacy bounds the information leakage about individual training records through the gradient updates, it does not prevent a compromised coordination server from using the gradients themselves as a channel for model extraction attacks — attempting to infer properties of participants' training data distributions through systematic gradient inspection. The threshold homomorphic encryption prevents this attack vector entirely by ensuring that the coordination server processes only encrypted gradients whose plaintexts it cannot access. The two security mechanisms are therefore complementary: differential privacy bounds per-record privacy leakage, and threshold homomorphic encryption prevents the coordination server from exploiting gradient information even in aggregate. Future work should formally verify that the composition of these two mechanisms provides the expected combined privacy guarantees under the composition theorems of differential privacy, accounting for the specific parameter choices used in the FedCRM production deployment.

The FedCRM evaluation confirms the practical applicability of federated learning to enterprise CRM analytics with heterogeneous data distributions, providing evidence that the theoretical promise of federated learning translates to real-world performance gains in the enterprise SaaS context. The research programme's multi-year investment in Salesforce security and analytics frameworks — LTDF, URGF,

CI/CD automation, HADES, GRAPHSEC, FedCRM — demonstrates a coherent progression from foundational security monitoring through governance, deployment security, AI quality assurance, supply chain security, and privacy-preserving analytics. FedCRM adds the privacy dimension to this programme: ensuring that the analytical capabilities built through this progression are deployable in contexts where data sharing constraints would otherwise prevent their use, enabling the full value of CRM analytics to be realised even in regulated, privacy-sensitive enterprise environments.

REFERENCES:

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, Apr. 2017, doi: 10.48550/arXiv.1602.05629. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html.

[2] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50–60, May 2020. [Online]. Available: doi: 10.1109/MSP.2020.2975749.

[3] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in Proc. Theory of Cryptography Conf. (TCC), 2006. [Online]. Available: doi: 10.1007/11681878_14.

[4] W. Verbeke, K. Dejaeger, D. Martens, J. Hur, and B. Baesens, "New insights into churn prediction in the telecommunication sector: A profit driven data mining approach," Eur. J. Oper. Res., vol. 218, no. 1, pp. 211–229, Apr. 2012. [Online]. Available: doi: 10.1016/j.ejor.2011.09.031.

[5] L. C. Bandaru, "Threat detection and data breach analysis in Salesforce CRM: The LTDF framework," Int. J. Innov. Res. Creative Technol. (IJIRCT), ISSN 2454-5988, vol. 7, no. 3, Jun. 2021. [Online]. Available: doi: 10.62970/IJIRCT.v7.i3.2605034.

[6] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in Proc. ACM CCS, 2017. [Online]. Available: doi: 10.1145/3133956.3133982.

[7] M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM CCS, 2016. [Online]. Available: doi: 10.1145/2976749.2978318.

[8] P. C. Jakku, L. C. Bandaru, and M. S. Bandrevu, "Automated vulnerability management in DevSecOps pipelines for SaaS platforms: A practical framework for SAST, DAST, dependency scanning, and controlled remediation," J. Adv. Dev. Res. (IJAIDR), E-ISSN 0976-4844, vol. 15, no. 1, Jan.–Jun. 2024. [Online]. Available: doi: 10.71097/IJAIDR.v15.i1.1904.

[9] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in Proc. AISTATS, 2020. [Online]. Available: https://proceedings.mlr.press/v108/bagdasaryan20a.html

[10] Salesforce, Inc., "Apex developer guide: Batch Apex," Salesforce Developer Docs, Apr. 2024. [Online]. Available: https://developer.salesforce.com/docs/atlas.en-us.apexcode.meta/apexcode/apex_batch.htm

[11] European Parliament and Council, "Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (GDPR)," Off. J. Eur. Union, vol. L 119, pp. 1–88, May 2018. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj